

Enhanced Stopping of Macro-Particles in Particle-in-Cell Simulations

J. May and J. Tonge

*Department of Physics & Astronomy, University of California, Los Angeles,
CA 90095*

W. B. Mori

*Department of Electrical Engineering, University of California,
Los Angeles and*

*Department of Physics & Astronomy, University of California, Los Angeles,
CA 90095*

F. Fiuza

Lawrence Livermore National Laboratory, CA 94550

R. A. Fonseca and L. O. Silva

*GoLP/Instituto de Plasma e Fusão Nuclear, 1049-001 Lisboa,
Portugal*

C. Ren

University of Rochester

(Dated: 24 January 2014)

We derive an equation for energy transfer from relativistic charged particles to a cold background plasma appropriate for finite-size particles that are used in particle-in-cell simulation codes. Expressions for one-, two-, and three-dimensional particles are presented, with special attention given to the two-dimensional case. This energy transfer is due to the electric field of the wake set up in the background plasma by the relativistic particle. The enhanced stopping is dependent on the q^2/m , where q is the charge and m is the mass of the relativistic particle, and therefore simulation macro-particles with large charge but identical q/m will stop more rapidly. The stopping power also depends on the effective particle shape of the macro-particle. These conclusions are verified in particle-in-cell simulations. We present 2D simulations of test particles, relaxation of high-energy tails, and integrated fast ignition simulations showing that the enhanced drag on macro-particles may adversely affect the results of these simulations in a wide range of high-energy density plasma scenarios. We also describe a particle splitting algorithm which can potentially overcome this problem and show its effect in controlling the stopping of macro-particles.

PACS numbers: 52.57.Kk, 52.38.-r

I. INTRODUCTION

Particle-In-Cell (PIC) simulations^{1,2} are a useful tool for studying nonlinear and kinetic physics such as those found in laser-solid interactions^{3,4}, fast ignition inertial fusion energy⁵⁻⁹, and relativistic collisionless shocks^{10,11}. Since PIC codes are based on first principles they are a preferred tool to explore physics in regimes where details of the distribution function affect the overall behavior of the system, where there have not been many experiments, or where experimental results are not well understood. For instance, in laser-solid plasma interactions, fast ignition, and relativistic shock studies fully kinetic simulations that resolve the electron dynamics are necessary. In these problems the flow of relativistic electrons in a background plasma is an important process. However, the physical scale length that needs to be resolved in these systems is the collisionless skin depth and not the Debye length. As the Debye length needs to be resolved in order to avoid numerical grid heating which results from aliasing, the resolution used in such simulations is sometimes much finer or the electron temperature is much higher than necessary. The use of higher order splines, and current smoothing and compensation can eliminate grid heating permitting larger cells and/or lower temperatures to be used^{2,12}. The use of larger cells also permits larger simulations that model the full spatial domains of the problems of interest. In order to carry out such large simulations in multiple dimensions it is common to limit the number of particles per cell, which is generally chosen at a value that keeps spurious finite size particle collisions and noise to a sufficiently small value.

Here we describe and quantify another effect that needs to be considered in PIC simulations when relativistic electrons are present^{13,14}. This work was originally motivated by studies of fast ignition relevant plasmas⁷. In these simulations a high intensity laser self-consistently generates a distribution of relativistic electrons with densities up to $\sim 10^{22}\text{cm}^{-3}$ moving through a $10^{23}\text{cm}^{-3} - 10^{24}\text{cm}^{-3}$ density background plasma^{6,7,15}. In a previous paper⁷ we observed an anomalous relaxation of the tail in the distribution function of these high-energy electrons in 10^{23}cm^{-3} density plasma; we ascribed this relaxation to the effect of plasmon emission, or wakes, in a turbulent plasma. Although plasmon emission is a physical effect, in this paper we show that its importance on a single particle is enhanced by the use of macro-particles in PIC codes. We show that the stopping power of relativistic electrons due to plasmon emission scales as q^2/m , where q is the charge and m is the mass of

the simulated relativistic electrons. In PIC simulations the ratio q/m of the simulated particles is kept consistent with real electrons, but the mass and hence the charge can be many times greater, thus strongly affecting the energy exchange between relativistic electrons and background plasma. In addition, the stopping power will depend on the dimensionality of the simulation and on the shape of the macro-particles. This artificially high plasmon emission can be controlled by using more particles per cell, larger cells, and higher order particle shapes. This is similar to how collisions and fluctuations are controlled in a PIC code. In this paper, we derive expressions for the stopping power of relativistic electrons in PIC simulations, and discuss how parameters can be chosen to reduce this effect with an emphasis on two-dimensional (2D) simulations of fast ignition relevant plasmas. These results are also relevant to general laser-solid interaction and relativistic shock simulations.

This paper is organized as follows. In Section II the equations for the macro-particle stopping of high energy particles in 2D PIC simulations by energy transfer through wakes are derived. The 1D and 3D cases are also discussed. In Section III the importance of macro-particle stopping in fast ignition simulations is discussed and the stopping of single particles in a nominal fast ignition plasma is analyzed. Section IV addresses the relaxation of high-energy tails of the electron distribution function due to macro-particle stopping. In Section V we turn to full-scale fast ignition simulations and analyze how varying the number of particles per skin depth in isolated target simulations changes the results. Section VI outlines a particle splitting technique we have developed to reduce the effect of stopping by reducing the size of high energy electrons. Finally, in Section VII we summarize our results.

II. THEORY

A charged particle moving faster than the thermal velocity in a plasma forms density wakes. The electric field of the wake at the location of the particle will do work on the particle and slow it down. To find the energy loss due to these wakes, we will consider a highly relativistic point electron moving through a cold fluid background plasma, and will solve for the electric field of the wake, following the methods used in the plasma-based accelerator community^{16,17}. We present results for the wake and electric field on the particle for 1D, 2D, and 3D. We then consider the effect of finite sized particles on the stopping, with an emphasis on 2D as it is presently the most common case.

We start by considering relativistic particles with charge ρ_b moving near the speed of light (c) in a cold fluid plasma, following the approach in Ref.¹⁶. The linearized fluid equations for the background plasma are:

$$\frac{\partial \vec{v}_1}{\partial t} = -\frac{e}{m_e} \vec{E}, \quad (1)$$

$$\frac{\partial n_1}{\partial t} = -n_0 \vec{\nabla} \cdot \vec{v}_1, \quad (2)$$

$$\vec{\nabla} \cdot \vec{E} = -4\pi e n_1 + 4\pi q n_b, \quad (3)$$

where e is the charge of an electron, m_e is its mass, v is the electron fluid velocity, E is the electric field, n is the plasma density, the subscripts 0 and 1 indicate zeroth and first order quantities respectively, and the subscript b corresponds to quantities for the relativistic particle species. We next consider the response for a single beam particle moving in the z direction with a speed v_b , in which case

$$q n_b = q \delta(\vec{r} - \vec{v}_b t) = q \delta(z - v_b t) \delta(\vec{r}_\perp) \quad (4)$$

where $\delta()$ is the Dirac delta function, q is the magnitude of the charge on the test particle, \vec{r} the position vector, and \vec{r}_\perp is the part of the position vector perpendicular to the direction of motion. The response to this point particle can also be viewed as the Green's function response of the electric field when one is calculating the response to a continuous beam distribution. We will also use Faraday's law

$$-\vec{\nabla} \times \vec{E} = \frac{1}{c} \frac{\partial \vec{B}}{\partial t}, \quad (5)$$

and Ampere's law where we have substituted for the current

$$\vec{\nabla} \times \vec{B} = -\frac{4\pi e n_0}{c} \vec{v}_1 + \frac{4\pi q n_b}{c} \vec{v}_b + \frac{1}{c} \frac{\partial \vec{E}}{\partial t}. \quad (6)$$

We combine these equations to get

$$\frac{\partial^2 n_1}{\partial t^2} + \omega_p^2 n_1 = \omega_p^2 \frac{q}{e} n_b = \frac{\omega_p^2 q}{v_b e} \delta(\vec{r}_\perp) \delta(t - \frac{z}{v_b}). \quad (7)$$

where $\omega_p = \sqrt{4\pi n_0 e^2 / m_e}$ is the unperturbed plasma frequency.

Due to causality the density perturbation in front of the relativistically moving test particle must vanish, so the density response is

$$n_1 = \frac{\omega_p q}{v_b e} \delta(\vec{r}_\perp) \eta(t - \frac{z}{v_b}) \sin(\omega_p(t - \frac{z}{v_b})) \quad (8)$$

where $\eta()$ is the Heaviside step function. To find the electric field, we use equations (1),(5), and (6) to derive a wave equation for the electric field,

$$-\frac{1}{c^2} \frac{\partial^2 \vec{E}}{\partial t^2} + \nabla^2 \vec{E} - \vec{\nabla} \vec{\nabla} \cdot \vec{E} = k_p^2 \vec{E} + \frac{4\pi}{c^2} \frac{\partial \vec{v}_b q n_b}{\partial t}. \quad (9)$$

where $k_p = \frac{\omega_p}{c}$ is the plasma wave number.

Next, we substitute from Gauss's law (3) and assume the test particle's velocity remains very close to c , i.e. $\vec{v}_b = c\hat{z}$, where \hat{z} is the unit vector in the z direction. We concentrate on the component of the electric field in the \hat{z} direction, and use the fact that both E_z and ρ_b are therefore functions of $(z - ct)$ such that $\frac{\partial}{\partial t} + c\frac{\partial}{\partial z} = 0$, to obtain

$$(\nabla_{\perp}^2 - k_p^2)E_z = -4\pi k_p q \frac{\partial}{\partial z} [\delta(r_{\perp}) \eta(t - \frac{z}{c}) \sin(\omega_p(t - \frac{z}{c}))] \quad (10)$$

The solution for E_z can then be written as

$$E_z = G_R(r_{\perp}) 4\pi k_p^2 q \eta(t - \frac{z}{c}) \cos(\omega_p(t - \frac{z}{c})) \quad (11)$$

where

$$(\nabla_{\perp}^2 - k_p^2)G_R(r_{\perp}) = \delta(r_{\perp}). \quad (12)$$

The solutions to Eq. (12) depend on the dimensionality of the problem. In 1D, the $\delta(r_{\perp})$ in Eq. (4) should be replaced by 1, ∇_{\perp}^2 vanishes, and q is in units of charge per unit area, so that $G_{R1D} = -1/k_p^2$. This yields an expression for the electric field

$$E_{z1D}(z) = -4\pi \bar{q} \eta(t - \frac{z}{c}) \cos(\omega_p(t - \frac{z}{c})) \quad (13)$$

where \bar{q} is the charge per unit area.

In 3D, the solution to Eq. (12) can be shown to be the Modified Bessel Function of the Second Kind¹⁸

$$G_{R3D}(r_{\perp}) = -\frac{1}{2\pi} K_0(k_p |r_{\perp}|) \quad (14)$$

which yields the Green's function for the electric field behind a test charge

$$E_{z3D}(r_{\perp}, z) = -2q k_p^2 K_0(k_p r_{\perp}) \eta(t - \frac{z}{c}) \cos(\omega_p(t - \frac{z}{c})), \quad (15)$$

The stopping power on a charge is the Green's function evaluated at the particle. For point particles the value of the step function is $1/2$ at the position of the particle. From Eq. (15), in 3D the electric field diverges logarithmically at the origin, which violates the

linearity condition. The stopping power would also diverge logarithmically. Therefore, if this response is viewed as the electric field from a real point electron then the result cannot be valid near the origin. However, when viewed from the stand point of a finite size particle (or beam), what matters is the solution for E_z obtained by integrating the Green's function over the particle's shape. As long as this electric field is finite then the use of linear theory can still be valid. We next carry out this integral in detail for the 2D case and comment on the 1D and 3D cases at the end of the section.

In 2D (slab) geometry, where $\nabla_{\perp}^2 \rightarrow \frac{d^2}{dr_{\perp}^2}$ and $\delta(r_{\perp}) = \delta(r_{\perp})$ it is readily verified¹⁷ that the solution to Eq. (12) is given by setting $G_{R2D}(r_{\perp}) = -\frac{1}{2k_p}e^{k_p r_{\perp}}$ for $r_{\perp} < 0$ and $G_{R2D}(r_{\perp}) = -\frac{1}{2k_p}e^{-k_p r_{\perp}}$ for $r_{\perp} > 0$. Thus the electric field is

$$E_{z2D}(r_{\perp}, z) = -2\pi e^{-k_p |r_{\perp}|} k_p \bar{q} \eta(t - \frac{z}{c}) \cos(\omega_p(t - \frac{z}{c})), \quad (16)$$

where \bar{q} is the charge per unit length. Unlike the 3D case, this electric field remains finite at the origin allowing direct analysis without specifying a particle shape. Using $1/2$ for the value of $\eta(0)$, the energy loss for a point (line of charge in 2D) particle is given by

$$\frac{d\epsilon_{2D}}{dt} = -\pi \omega_p \bar{q}^2 \quad (17)$$

where ϵ is the energy of the test particle. Dividing by $\omega_p \bar{m} c^2$ we get

$$\frac{d\gamma_{2D}}{d\omega_p t} = -\pi \frac{\bar{q}^2}{\bar{m} c^2}. \quad (18)$$

where γ is the relativistic Lorentz factor of the test charge. Equation (18) shows that the stopping power (distance) for a relativistic particle in a cold plasma is (inversely) proportional to the \bar{q} on the particle holding the \bar{q}/\bar{m} ratio fixed. To derive an equation for PIC simulations we use the fact that $\bar{q}/e = n_0 \Delta^2 / N$ where Δ is the cell size (we assume equal cell sizes along the different dimensions) and N is the number of simulation macro-particles per cell, and obtain

$$\frac{d\gamma_{2D}}{d\omega_p t} = -\frac{1}{4} \frac{\omega_p^2}{c^2} \frac{\Delta^2}{N} C. \quad (19)$$

Equation (19) demonstrates that for large cell sizes N must be increased to reduce the macro-particle stopping to a desired level. The factor C is a term that accounts for finite size particle effects which will depend on the details of the particle shape and smoothing used in the simulation.

Next we consider finite size particle effects and the factor C . These effects will reduce the stopping power of the test particle. The use of finite size particles reduces both the electric field made by the test charge and the force from this electric field on the particle. We first discuss the reduction of the electric field. The effect of the finite size of the particle on the electric field is found from a convolution of the point particle electric field with the shape factor as a function of grid size

$$\bar{E}_z(r_\perp, \xi) = \int_{-\infty}^{\infty} dr'_\perp \int_{-\infty}^{\infty} d\xi' E_z(r_\perp - r'_\perp, \xi - \xi') S(r'_\perp, \xi'), \quad (20)$$

with $\xi = (z - ct)$. For linear splines (area weighting) the shape factor is given by

$$S_1(x) = \frac{1}{\Delta} \left(\left(1 - \frac{x}{\Delta}\right) \eta(x) \eta(\Delta - x) + \left(1 + \frac{x}{\Delta}\right) \eta(-x) \eta(\Delta + x) \right), \quad (21)$$

where x is the particle position and Δ is the cell size in a specified direction. As noted above we consider square cells in 2D but the arguments presented can be easily extended to rectangular cells. The Fourier transform of this shape function, which is an order 1 spline, is a sinc function squared. Higher order spline shapes are defined as the convolution of the previous order with the zeroth order shape (nearest grid interpolation),

$$S_n(x) = \int_{-\infty}^{\infty} dx' S_0(x - x') S_{n-1}(x'). \quad (22)$$

Therefore, for splines of order n the Fourier transform of the shape function is given by

$$S(k)_n = \left(\frac{\sin(\frac{k\Delta}{2})}{\frac{k\Delta}{2}} \right)^{n+1}. \quad (23)$$

The shape factor is dependent on the cell size and on the interpolation order. In addition, digital filters are also commonly applied. A filter of particular interest used for current smoothing is a triangular (1-2-1) filter applied multiple times in each direction. A compensator² is often used to flatten the profile in k -space for low k . These filters take the k -space form

$$F(k)_n = \cos^{2n} \left(\frac{k\Delta}{2} \right) \quad (24)$$

where here n is the number passes of the triangular filter. When the compensator is included the effective filter takes on a more complex expression, but there is no k dependence in the denominator so smoothing (including compensation) is never as effective as using higher-order splines at reducing aliased noise (the source of particle heating in PIC codes).

For illustrative purposes, we examine the electric field generated by different particle shapes where the charge on the particle is kept fixed. Figure 1 shows the magnitude of the axial (z) component of the electric field produced on axis ($r_{\perp} = 0$) as a function of z by a 2D relativistic particle of finite size (square) in a plasma according to equation (20). In figure 1a the particle shape is a second order spline, i.e., $n = 2$ in Eq. (23), and the cell size is varied from a point particle, to $0.5c/\omega_p$, to $1.0c/\omega_p$, and to $2.0c/\omega_p$ (blue, red, green and magenta curves respectively). These curves confirm that increasing the cell size for a fixed charge reduces the electric field strength and smooths it out over the location of the particle. In Figure 1b, we show that increasing the interpolation order while keeping the cell size and charge on the particle fixed also decreases the amplitude of the electric field. We plot the electric field for linear (blue), quadratic (magenta), cubic (green), and quartic (red) interpolation while holding cell size fixed at $2.0c/\omega_p$. This shows that each higher order spline corresponds to a ‘wider’ particle. Varying the cell size or varying the particle order both effectively change the size of the particle, however, the use of larger cells also changes the accuracy of the field solver for how light waves propagating in vacuum. We also note that although the macro-particle stopping effects are reduced as the cell size is increased to values larger than a skin depth (and Debye length), other issues such as significant modifications of the dispersion properties of both plasma and electromagnetic waves will become important.

Next, we turn our attention from the electric field of the wake to the stopping power of the relativistic particles. The stopping power is obtained from an integral of the electric field in Eq. (20) with the shape function

$$F_{stopping} = q \int_{-\infty}^{\infty} dr'_{\perp} \int_{-\infty}^{\infty} d\xi' \bar{E}_z(\vec{r}'_{\perp}, \xi') S(\vec{r}'_{\perp}, \xi'). \quad (25)$$

It is worth noting that total force on the particle can also be obtained by integrating the shape function and electric field in k space,

$$F_{stopping} = q \int_{-\infty}^{\infty} d\vec{k} S^2(\vec{k}) E_z(\vec{k}), \quad (26)$$

where through Eq. (21) it can be seen how the use of progressively higher order splines (and smoothing) can reduce the stopping force. The shape functions can be different in each direction, but it is generally a separable function in the each coordinate in which only the cell size varies. For separable shape functions, it is easy to extend eqs. 25 and 26 to multi-dimensions. For simplicity, in what follows we assume the cell sizes are squares in 2D.

It is difficult to obtain analytical expressions for $F_{stopping}$ for the particle shapes used in PIC codes, therefore we simply plot numerical solutions to eq. (25). Figure 2 shows the effect of varying the spline order and current smoothing on the stopping force. We plot the force on the particle for linear and cubic splines, and the same using 4-pass smoothing (which we often use in Osiris) and k-filtering (used in Parsec); it is normalized to the stopping for a point particle, and therefore this curve is equal to the factor C for two dimensions, C_2 , in Eq. (19). As stated earlier, both smoothing and higher-order interpolation schemes reduce the macro-particle stopping. However, for fixed charge per particle, the reduction is only significant for large cell sizes ($\Delta > 1/k_p$), where important skin-depth physics can be missed and where significant numerical dispersion effects can be present². In essence by increasing the cell size, one is smoothing out the wake field, whose scale length is on order of the skin depth. For example, for a cell size of $1/2 k_p^{-1}$, without current smoothing, the stopping power is 77% and 68% of the point-particle limit for linear and cubic interpolation respectively. If the cell size is increased to $3/2 k_p^{-1}$, then these values become 37% and 21% respectively. Including current smoothing further reduces the value of C to 0.09 in the linear case and 0.05 in the cubic one. We note that there are numerical limits to both increasing cell size and increasing interpolation order for the purpose of reducing macro-particle stopping. For example, the use of large cells can lead to plasma waves and light waves having negative group velocity. In finite-difference PIC codes where numerical errors cause the phase velocity of electromagnetic waves on the grid to be less than the speed of light, relativistic particles can excite numerical Cherenkov noise which is increased as the time step is reduced below the electromagnetic Courant condition for stability². In addition, since the hard stability limit for the time step in a plasma is $2/\omega_p^2$ this introduces a practical limit on the cell size of $2\sqrt{2}c/\omega_p$ when dealing with relativistic particles. As seen in equation (19) the stopping power can be reduced by also increasing the number of particles per cell (or effectively decreasing the particle charge). This is the best method with respect to not modifying the physics, but it is obviously the most computationally costly method.

We close this section with a brief discussion on macro-particle stopping in 1D and 3D. The electric field for an infinitesimally thin sheet of charge was given in Eq. (13). We note that it could also be obtained from integrating the 3D result, Eq. (15), over a beam with a charge density that is uniform across the beam in the transverse direction. Interestingly the 1D result is independent of the plasma density, however, the result is only valid if the electric

field is still sufficiently small for linear theory to be valid. This is true if the normalized field value is small, $\frac{eE_z}{mc\omega_p} \ll 1^{16,17}$. For a macro-particle one needs to integrate Eq. (13) over the shape of the particle. The energy loss for a point (sheet of charge in 1D) particle is given by

$$\frac{d\epsilon_{1D}}{dt} = -2\pi\bar{q}^2c, \quad (27)$$

which can be written for the Lorentz factor of the test charge as

$$\frac{d\gamma_{1D}}{d\omega_pt} = -2\pi\frac{\bar{q}^2}{\bar{m}\omega_pc}. \quad (28)$$

Taking into account the fact that $\bar{q}/e = n_0\Delta/N$ in 1D, we obtain the energy loss in 1D PIC simulations

$$\frac{d\gamma_{1D}}{d\omega_pt} = -\frac{1}{2}\frac{\omega_p}{c}\frac{\Delta}{N}C_1 \quad (29)$$

where C_1 is the effect of the particle shape in 1D.

In 3D the Green's function logarithmically diverges as $|\vec{r}_\perp| \equiv r$ approaches zero. However, for a finite size particle the divergence disappears. A detailed analysis is beyond the scope of this work. However, in¹⁹ the result for flat top and Gaussian particle (beam) shapes in the transverse coordinate and Gaussian shapes in the longitudinal directions were given. Here we summarize the results for Gaussian shaped macro-particles, for which the charge is given by $\frac{q}{(2\pi)^{3/2}}e^{-(z-ct)^2/2\sigma_z^2}e^{-r^2/2\sigma_r^2}$, where σ_z and σ_r are the particle size in the longitudinal and transverse directions respectively. For such a charge the amplitude of the wakefield is given in¹⁹ as

$$E_{z3D} = -qk_p^2e^{-k_p^2\sigma_z^2/2}e^{k_p^2\sigma_r^2/2}\Gamma\left(0, \frac{k_p^2\sigma_r^2}{2}\right), \quad (30)$$

where $\Gamma(\alpha, \beta) = \int_\beta^\infty s^{\alpha-1}e^{-s}ds$ is the incomplete gamma function. For symmetric macro-particles, $\sigma_r = \sigma_z = \sigma$, we have simply

$$E_{z3D} = -qk_p^2\Gamma\left(0, \frac{k_p^2\sigma^2}{2}\right). \quad (31)$$

We note that in the limit that $k_p\sigma$ approaches zero, $\Gamma(0, \frac{k_p^2\sigma^2}{2}) \approx \ln(1.12/k_p\sigma)$, and the electric field on axis reduces to

$$E_{z3D} = -qk_p^2\ln\left(\frac{1.12}{k_p\sigma}\right). \quad (32)$$

This is the peak amplitude of the wake. To derive how the electric field is distributed over the particle requires carrying out the full integrals over the particle which is beyond the

scope of this work. However, this expression is useful for estimating the macro-particle stopping and we will use it in what follows.

Following the same procedure as in the 1D and 2D cases, the energy loss in 3D is then described by

$$\frac{d\epsilon_{3D}}{dt} = -q^2 \frac{\omega_p^2}{c} \ln \left(\frac{1.12}{k_p \sigma} \right), \quad (33)$$

or, for the Lorentz factor of the test charge, as

$$\frac{d\gamma_{3D}}{d\omega_p t} = -\frac{q^2}{m} \frac{\omega_p}{c^3} \ln \left(\frac{1.12}{k_p \sigma} \right). \quad (34)$$

If we assume cubic cell shapes with the same cell size, Δ in each direction then $q/e = n_0 \Delta^3 / N$ and we the cell size for the particle size, σ , then we can estimate the macro-particle stopping in 3D, which is given by

$$\frac{d\gamma_{3D}}{d\omega_p t} = -\frac{1}{4\pi} \frac{\omega_p^3}{c^3} \frac{\Delta^3}{N} \ln \left(\frac{1.12}{k_p \Delta} \right) C_3. \quad (35)$$

where C_3 is a factor that accounts for additional macro-particle effects. We note that for a fixed number of particles per cell, N , and a cell size $\Delta \leq c/\omega_p$ (neglecting the exact details of the particle shape), the macro-particle stopping decreases with increasing dimensionality. Therefore, it requires a larger number of particles per cell in 1D and 2D to control this effect when compared with 3D. We note that the formalism that has been presented can be used to carry out an exact and detailed analysis of the 3D case as we did for the 2D case.

III. SINGLE PARTICLE STOPPING IN FAST IGNITION SIMULATIONS

We can now make predictions for the stopping of relativistic macro-particles in background plasmas. Here we focus on the stopping distance in two dimensional fast ignition relevant simulations. To motivate the importance of macro-particle stopping for fast ignition studies, let us take as example the parameters used in integrated PIC simulations of fast ignition⁶⁻⁸ aimed at modeling both laser absorption and transport of fast electrons in a background plasma. We compute the stopping power for the simulations with different resolutions and numbers of particles per cell. In these simulations an intense laser with a 1 μm wavelength interacts with a plasma with a nominal density around $100n_c$ (where n_c is the critical plasma density for the laser). The number of macro-particles per skin depth squared (PPSD²) in these simulations varied between 6 and 24, the cell sizes varied from 0.5

to $2 k_p^{-1}$, and the particle order varied from linear to cubic. Despite these differences, the predicted energy loss rate was similar and it only varied in the range $0.12 - 0.34 \text{ MeV}/\mu\text{m}$. The simulations used different plasma lengths so the single particle energy loss across the simulation box varied between 4 and 18 MeV. For example, in the paper by Tonge et al.⁷ the cell size was $0.5 c/\omega_p$, with 4 particles per cell ($N = 4$) giving 16 PPSD², and quadratic splines and 5-pass compensated smoothing were used, leading to a stopping force of $0.011 m_e c \omega_p$, or an energy loss rate of $0.337 \text{ MeV}/\mu\text{m}$ at a density of $100n_c$. In each of these cases the energy loss of a macro-particle due to enhanced stopping is comparable or larger than the typical fast-electron energies relevant for fast ignition. Therefore the enhanced stopping of macro-particles could greatly affect the results in each case.

We verify these predictions by running simulations of single particle stopping with the PIC codes OSIRIS²⁰⁻²² and PARSEC²³. The first is a finite difference code, whereas the second is a spectral code which will not produce numerical Cherenkov radiation of a single particle without the effect of aliasing (note that there is radiation from aliasing effects). In OSIRIS we also use current smoothing and compensation, while in PARSEC the particles also have a Gaussian shape with a size equal to the grid size. For the single particle stopping tests, we run 2D simulations with periodic boundary conditions. In these simulations the cell size is $0.5 c/\omega_p$, the box size is $128 c/\omega_p \times 128 c/\omega_p$ and quadratic splines are used for interpolation. Both ions and electrons are included, the mass ratio being 3672 (Deuterium plasma). Simulations were performed with $N = 64$ in the background plasma, the background temperature is 100eV, and the duration of the simulations is $44.8 \omega_p^{-1}$. In order to study the wake driven by a single relativistic electron in the background plasma we use a subtraction technique^{24,25}, where two identical simulations are run, but in one of them a single 50 MeV electron moving in the z direction (hereafter, x_1 direction) is added. The wake electric field is obtained by subtracting the electric fields of both simulations. This tracer particle has the same \bar{q} as particle with 16 particles per $1/k_p^2$.

Figure 3a shows the E_1 obtained using the subtraction technique from PARSEC simulations, thereby eliminating the noise and clearly showing the wake of the fast particle. In figure 3b we show a zoomed-in line-out on axis of the wake field from 3a (red), along with the theoretical prediction from equation (15) (blue). As can be seen the results are in good agreement, although the simulation does produce a more ragged and slightly weaker field than predicted by theory, even with the finite sampling taken into account.

We also compared the energy loss observed in the simulations to the theory presented in Section II. According to equation Eq. (19) a point particle will lose an energy of $0.7 m_e c^2$ in $44.8 \omega_p^{-1}$; including finite-size effects this is decreased to $0.30 m_e c^2$ in PARSEC and $0.47 m_e c^2$ in OSIRIS. In our simulations the tracer particle lost $0.44 m_e c^2$ in PARSEC and $0.64 m_e c^2$ in OSIRIS. Both results are on the order of and less than the point particle result, and the ratio of the two values is correct to theory. That both results show weaker finite-particle effects than predicted by macro-particle theory may be due to the background plasma being discrete and warm. Understanding these differences is an area for future work.

We have also performed a series of test particle simulations with PARSEC to examine the effect of background temperature and particle energy. Recall that the theory assumes a cold fluid plasma and that the particle is highly relativistic. These 2D PIC simulations use the same simulation box and numerical parameters of the previous simulations, but we varied the energy of the relativistic electron (> 1 MeV) and of the background plasma (100 eV - 1MeV). For background temperatures up to 10keV and all fast electron energies examined, the energy loss is relatively insensitive to either parameter, and is $\sim 14\%$ less than the theoretical value. For background temperatures > 100 keV, the velocity of background particles becomes relativistic, the assumption of a cold background plasma does not hold and the simulation results deviate more strongly.

We have tested $N = 4, 16, 64$, and 256 for the background plasma and the differences to the stopping of a single particle were negligible for background temperatures below 10keV. We therefore conclude that equation (19) is useful for all initial energies above 1MeV and for background temperatures up to ~ 10 keV.

IV. RELAXATION OF HIGH-ENERGY TAILS

We now investigate the macro-particle stopping for a distribution of fast electrons with a high-energy tail. These simulations are run with PARSEC and model a spatially uniform cold (1 keV) background plasma with a relativistic tail distribution with a slope temperature of 2.5 MeV and with 4% of the particles in the tail. The simulations have periodic boundaries, with a box size of 256 cells by 256 cells, a cell size of $0.5c/\omega_p$, and $N = 16$ or 64 PPSD². Second order particle interpolation is used with Gaussian shaped particles with particle sizes equal to the grid size. In Figure 4 the blue curve shows the initial energy distribution, with

the red, green and yellow curves showing the relaxation of the distribution with time, 92.75 fs apart when scaled to a background density of 10^{23} cm^{-3} . The curves show that the high-energy particles uniformly lose energy as predicted by theory. The predicted energy loss is $0.09 \text{ MeV}/\mu\text{m}$ ($0.027 \text{ MeV}/\text{fs}$) and the simulations show an energy loss of $0.06 \text{ MeV}/\mu\text{m}$, which is less than both theory and single particle simulations.

In single particle simulations the energy loss is 86% of theory and insensitive to background particle count; in simulations with a high-energy tail the discrepancy with theory is dependent on the PPSD² because this affects the distribution of particles in the tail. For simulations where we vary the PPSD² but keep other parameters constant, we found energy loss of 71%, 66%, and 53% of the predicted value for 16, 64, and 256 PPSD² respectively. This effect is likely due to the more even distribution of particles in the tail across cell boundaries as N (not necessarily PPSD²) increases. For the particles in the tail the number of particles per cell varied from $N = 0.16, 0.64$, and 2.56 respectively (recall that 4% of the particles are in the tail). As the tail becomes more uniformly distributed in space the discrete effect of macro-particles is decreased. In simulations relevant to fast ignition the electrons are accelerated by the laser in bunches⁹ separated by half a laser wavelength which is comparable to a plasma wake wavelength. Therefore, the particles in these bunches will not interact through individual wakes but rather through collective wakes. If the bunch duration is short compared to a skin depth the wakes will collectively add leading to enhanced stopping while if the particles are bunched on scales much longer than the skin depth then the stopping power will be comparable to the single particle stopping power.

V. INTEGRATED FAST IGNITION SIMULATIONS

We will now revisit the integrated simulations using isolated fast ignition targets from Tonge et al.⁷. In these simulations a $50\mu\text{m}$ radius $100n_c$ target with a $20\mu\text{m}$ core region is illuminated with an $I = 8 \times 10^{20} \text{ W}/\text{cm}^2$, $\lambda = 1\mu\text{m}$ laser with a $20\mu\text{m}$ spot size; detailed parameters are given in Tonge et al.⁷. Here we reproduce these simulations using cubic splines (quadratic splines were used in the original simulations) and varying the PPSD² while keeping all other parameters fixed. By increasing the PPSD² from 16 (original value) to 100 we decrease the stopping by 6.25 times and we observe a significant change in the dynamics of the simulation due to macro-particle stopping. The stopping decreases the

energy of hot electrons while increasing the heating of the background plasma between the laser-plasma interface and the target core. Figure 5 shows the momentum distribution function in the isolated target, 835 fs after the laser strikes, at the longitudinal (x_1) position $2\text{ }\mu\text{m}$ behind the laser-plasma interface and integrated along the transverse direction (x_2) for the 16 PPSD² or 100 PPSD² cases. The blue curve (16 PPSD²) is clearly wider than the red curve (100 PPSD²) indicating the plasma is hotter in the interaction region. Panel b) shows the power absorbed in the target core as a function of total laser power. From the initial heating at 400 fs up to 600 fs the two simulations have similar power delivered to the core, although the red curve has a bump at 450 fs which is due to higher level of refluxing off the back of the target early in the high particle count simulation. After 600 fs the curves diverge, with the lower PPSD simulation showing a larger heating of the core; this is due to both energy of the electrons reaching the core being lower and hence the stopping of fast electrons in the core being enhanced, and to the increased heating of the background plasma in front of the core that causes heated background electrons to travel into the core and redeposit their absorbed energy. This overall effect was to increase the power delivered to the core by 25% at 1.3 ps. In the lower absorption case (higher PPSD case) the laser is still capable of delivering 10% of its power to the core.

Let us now compare the magnitude single particle stopping effect to that seen by a real particle (although these are 2D simulations, the 3D result is equivalent up to a factor of $\mathcal{O}(1)$.) For a $1\mu\text{m}$ laser, $100n_c = 10^{23}\text{ cm}^{-3}$, the cell size of $0.05c/\omega_0$ is $\sim 0.008\mu\text{m}$, giving 5×10^4 electrons in a cell volume. For $N = 25$, and taking into consideration finite particle size effects, the stopping is ~ 1300 times stronger in the simulation than would be seen for a physical electron. To make the real and simulated macro-particle stopping powers numerically equivalent at $n = 100n_c$ and retain the skin depth physics (i.e. cell size $\leq c/\omega_p$) we would have to increase the particle count to $N > 500$ and use current smoothing and cubic splines. This is impractical for multi-dimensional PIC studies of relevant fast ignition scales. However, it is not important to make the macro-particle stopping power comparable to the real stopping power, rather it is important to make it small enough over the plasma size so that this enhanced stopping does not change the results. On the other hand, if one is modeling the heating of the core at very high densities, then since the stopping power of the electrons is the physics process that is most important and it needs to be quantitatively correct. This is similar to the effect of the collisions between macro particles. If one is

modeling a region of plasma for which collisions are not important it is not important that the collision frequency is correct, rather it must be kept small enough that the growth rates and dispersion properties of collisionless processes are not modified.

VI. PARTICLE SPLITTING

We now discuss a scheme that can potentially overcome this effect of macro-particle stopping in PIC simulations associated with laser-solid interactions, fast-ignition, and relativistic shocks. We have shown that reducing the charge on the relativistic, i.e., ‘beam’, electrons is an effective method to alleviate the enhanced stopping of relativistic particles due to their wakes. However, in the majority of the simulations of interest we want to study the way particles are accelerated, e.g. by the laser in fast ignition scenarios or by the shock in astrophysical plasmas, and we do not know beforehand which particles will be accelerated (fast) and which will be background plasma particles. Thus, in a standard PIC simulation reducing the charge of a fast particle requires reducing the initial charge of all simulation particles, leading to a prohibitive particle count.

In many of these scenarios, the super-thermal particles which are stopped by macro-particle effects are only a small percentage of the total particle count - although at the same time these particles can have a profound effect on plasma dynamics. Under these conditions an algorithm that splits the relativistic electrons into macro particle electrons with less charge might be effective. We have implemented and experimented with a splitting algorithm that works as follows. For a given simulation we define the minimum particle energy above which macro-particle stopping will become important and the maximum particle charge that high-energy particles should have for accurate description of their stopping. In the simplest algorithm these can be fixed values, for instance associated with the initial laser and plasma conditions. It would also be possible to make them dynamic values that depend on for instance the local background plasma conditions that vary due to density gradients or plasma steepening. The simulation is initialized in the standard way, however, after each n time steps (in the future the value of n could change statically or dynamically as the simulation progresses), we calculate the energy and charge of each particle and compare it to the minimum energy and maximum charge defined. If they are both larger, then we reduce its charge and mass by a factor of two and then duplicate the particle. This process is repeated

until the particle charge is smaller than the defined value. Thus, we effectively increase the particle count for the high-energy part of the distribution function. It is important to note that in most cases, because of the wide difference between the background thermal and ‘fast’ particle energies, the efficiency of the splitting is relatively insensitive to the exact energies chosen as the splitting points. Also, since the number of fast particles is significantly smaller than the number of background particles, the computational overhead of splitting is not dramatic. We also have found the results are relatively insensitive to the choice of n .

It is important to ensure that the duplicated particles move away from each other. If the duplicated particles have the exact same position and momentum then they will never move apart and their wakes will add coherent and the stopping power will remain unchanged. Furthermore, the use of finite size particles means that particles that start off close together within a cell will move apart slowly. Several schemes can be thought of to separate the duplicated particles, such as randomly shifting the position of the new particle within the same cell, slightly shifting the momentum of the two particles in a way that conserves the total momentum (this can introduce a small divergence), or simply relying on a collisional operator, which is already used in many of the scenarios of interest, to naturally separate the two particles in phase space. We have experimented both with the momentum shifting scheme and with the use of a Monte-Carlo binary Coulomb collisions operator^{12,26–28}, with overall satisfactory results.

We next show results using the OSIRIS Coulomb collisions module for electron-ion collisions. We initialize a 2D plasma in a box $1\mu\text{m}$ wide and periodic in the transverse direction, and $200\mu\text{m}$ long in the longitudinal direction. The box is filled with a $100n_c$ plasma, except for a thin ($2\mu\text{m}$) vacuum layer to the left which isolates the plasma from this wall; a laser with $I = 5 \times 10^{19} \text{ W/cm}^2$ (normalized vector potential $a_0 = 6$) and $\lambda = 1\mu\text{m}$ is incident from this direction. In order to avoid using a larger simulation box to capture the self-consistent plasma expansion due to laser heating of the plasma surface, we used infinitely heavy ions. We use a linear density ramp from 0-100 n_c in the first micron of plasma. We use cells with a size of $0.5c/\omega_p$, and either 16 or 512 particles per cell, corresponding to 64 or 2048 PPSPD², respectively; the time step satisfies the Courant condition almost exactly, to reduce numerical Cherenkov. We use third order particle shapes, with 4-pass current smoothing with a compensator.

In figure 6 we show the particle count and forward heat-flux of a typical laser-solid

simulation as a function of kinetic energy on a logarithmic scale for $N = 512$ (note that particle count is not corrected for bin size). The first thing to notice is that the majority of the energy flux is being carried by particles with $\gamma \sim 2 - 10$, demonstrating that the energy is being carried by relativistic electrons and that macro-particle stopping needs to be considered. It is also clear that the peak of the heat-flux and the peak of the particle count are separated by two orders of magnitude indicating that almost all of the forward heat flux is carried by only a few percent of the particles. This illustrates that if decreased the particle charge and increased the particle count of only a small region of phase space we could greatly minimize the importance of macro-particle stopping without significant computational expense. To illustrate the effectiveness of the splitting algorithm, in Figure 7 we show the total forward heat flux as a function of longitudinal position for the $N = 512$ case as well as a $N = 16$ case with no splitting and a $N = 16$ case with five particle splits (for final particle size equivalent to $N = 16 \times 2^5 = 512$). Figure 7b clearly shows significant loss of beam energy going into the plasma for the $N = 16$ case as compared to the $N = 512$ case. In fig. 7b the heatflux essentially vanishing by $x_1 = 100\mu\text{m}$. Taking the peak beam particle energy as $2a_0m_e c^{29}$, we find for these parameters a stopping distance of $97.5\mu\text{m}$, consistent with these results. For the $N = 512$ case (Fig. 7a), we predict a decrease of the particle Lorentz factor, γ , of 0.375 over the same distance, which we estimate to be 3 – 6% of beam particle energy. The larger apparent loss of heat flux moving forward in the case of $N = 512$ is due to a number of non-macro-particle effects, including the increase of laser absorption as a function of time combined with the time-of-flight of electrons, the fact that not all heat flux is being carried by highly-relativistic particles, and the effects of the wakes of at least partially coherent particle bunches. We note that collisional stopping, however, is not a significant factor in this regime. More importantly, the results using $N = 16$ but with 5 binary particle splits (Fig. 7c) are identical to the $N = 512$ simulation. Figure 7d shows the comparison of the heat flux carried in the forward direction as a function of the longitudinal positions for the different cases. Again, it is possible to observe that by splitting the high-energy particles the macro-particle stopping is controlled. For the example shown, using the splitting algorithm leads to only a factor of two increase in the computational time, giving a total computational savings factor of 16 to reach the same accuracy in terms of macro-particle stopping. This illustrates the efficiency of this algorithm and its usefulness for carrying out multi-dimensional studies of large plasma volumes. We should also note that

using this splitting algorithm can bring important statistical advantages when computing collisions between fast and background particles in many of the scenarios of interest. We note that that this algorithm can lead to load balancing issues. OSIRIS does have a dynamic load balancing capability and we will experiment with this as part of future work.

VII. SUMMARY

In this paper we have shown that relativistic particles moving in a cold background plasma in PIC simulations are susceptible to enhanced stopping due the use of macro-particles. The stopping scales as q^2/m so that particles with large charge but with the correct charge to mass ratio will stop more rapidly. This stopping is due to the wakefield created by relativistic particles and it can be predicted using wakefield theory developed for studying plasma wakefield acceleration. We reviewed the derivation of the wakefields created by a point particle moving near the speed of light for one dimension (charge sheet), two dimensions (line of charge), and three dimensions. We used this wakefield (Green's function) to derive the wakefields created by finite size particles (such as those used in PIC simulations), and also calculated the force on a particle from its own wakefield to get the stopping power. We also studied how the stopping depends on the cell size, particle shape, and dimensionality. The enhanced stopping is mitigated through the use of larger cells, higher order particle shapes, and current smoothing as well as with a decrease in the macro-particle charge. We found good agreement between the theory and results against PIC simulation from a finite difference PIC code (OSIRIS) and a spectral PIC code (PARSEC). We also studied how a distribution of electrons containing a hot tail relaxes and showed that the macro-particle stopping process dominates how the tail relaxes. We also reexamined previous results in intense laser-solid interactions, such as the isolated target simulations in⁷ and found that indeed this effect modified the physics. When the charge per particle was reduced by a factor of 6.25 the amount of laser energy being deposited in the core decreased by only 25% from the values quoted in⁷. Our expressions also predict that simulations done by others will also be impacted by this effect^{6,8}. Lastly, we described a particle splitting algorithm that can reduce this effect by decreasing the charge of high-energy particles as the simulation evolves. We found that this algorithm can successfully control the macro-particle stopping in PIC simulations of high-energy density scenarios in a computationally efficient way.

ACKNOWLEDGMENTS

We acknowledge a useful discussion with Dr. Max Tabak which was an initial motivation for this work in 2009. When finalizing this manuscript it came to our attention that independent work on this topic has been recently carried out by Kato²⁹. This work was supported by the US Department of Energy under the Fusion Science Center on Extreme States of Matter and Fast Ignition Physics through a subcontract from the University of Rochester, and by contracts de-na0001833, de-sc0008316, and de-sc0008491, and the NSF grant number NSF-ACI-1339893. This work was also performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. FF acknowledges financial support by the LLNL Lawrence Fellowship. The work of LOS and RAF is supported by the European Research Council (ERC-2010-AdG Grant No. 267841). The simulations were carried out on the Hoffman and Dawson2 Clusters at UCLA and on Intrepid at Argonne National Laboratory.

REFERENCES

- ¹J. M. Dawson, Rev. Mod. Phys. **55**, 403 (1983).
- ²C. Birdsall and A. Langdon, *Plasma Physics via Computer Simulation*, 1st ed. (Taylor and Francis, 1991).
- ³L. O. Silva, M. Marti, J. R. Davies, R. A. Fonseca, C. Ren, F. S. Tsung, and W. B. Mori, Phys. Rev. Lett. **92**, 015002 (2004).
- ⁴F. Fiuza, R. A. Fonseca, J. Tonge, W. B. Mori, and L. O. Silva, Phys. Rev. Lett. **108**, 235004 (2012).
- ⁵M. Tabak, J. Hammer, M. E. Glinsky, W. L. Kruer, S. C. Wilks, J. Woodworth, E. M. Campbell, M. D. Perry, and R. J. Mason, Physics of Plasmas **1**, 1626 (1994).
- ⁶B. Chrisman, Y. Sentoku, and A. Kemp, Physics of Plasmas **15** (2008).
- ⁷J. Tonge, J. May, W. B. Mori, F. Fiuza, S. F. Martins, R. A. Fonseca, and L. O. Silva, Physical Review E **16** (2009).
- ⁸A. Kemp, B. Cohen, and L. Divol, Physics of Plasmas **17** (2010), 10.1063/1.3312825.
- ⁹J. May, J. Tonge, F. Fiuza, R. A. Fonseca, L. O. Silva, C. Ren, and W. B. Mori, Phys. Rev. E **84**, 025401 (2011).

- ¹⁰A. Spitkovsky, The Astrophysical Journal Letters **673**, L39 (2008).
- ¹¹S. F. Martins, R. A. Fonseca, L. O. Silva, and W. B. Mori, The Astrophysical Journal Letters **695**, L189 (2009).
- ¹²Y. Sentoku and A. J. Kemp, J. Comp. Phys. **15**, 056309 (2008).
- ¹³J. Tonge, J. May, W. B. Mori, F. Fiuza, and L. O. Silva, in *Proceedings of the 40th Annual Anomalous Absorption Conference* (2010).
- ¹⁴A. J. Kemp, F. Fiuza, A. Debayle, T. Johzaki, W. B. Mori, P. K. Patel, Y. Sentoku, and L. O. Silva, arxiv:1308.2628 (2013).
- ¹⁵A. Kemp, Y. Sentoku, and M. Tabak, Physics of Plasmas **79**, 066406 (2009).
- ¹⁶T. Katsouleas, S. Wilks, P. Chen, J. M. Dawson, and J. J. Su, Particle Accelerators **22**, 81 (1987).
- ¹⁷W. B. Mori, J. M. Dawson, C. Joshi, T. Katsouleas, J. J. Su, and S. Wilks, “The plasma wakefield accelerator,” (1988).
- ¹⁸M. Abramowitz and I. Stegun, *Handbook of mathematical functions* (1970).
- ¹⁹W. Lu, C. Huang, M. M. Zhou, W. B. Mori, and T. Katsouleas, Physics of Plasmas **12**, 063101 (2005).
- ²⁰R. A. Fonseca, L. O. Silva, F. S. Tsung, V. K. Decyk, W. Lu, C. Ren, W. B. Mori, S. Deng, S. Lee, T. Katsouleas, and J. C. Adam, Lecture Notes in Computer Science **2331**, 342 (2002).
- ²¹R. A. Fonseca, S. F. Martins, L. O. Silva, J. W. Tonge, F. S. Tsung, and W. B. Mori, Plasma Physics and Controlled Fusion **50**, 124034 (9pp) (2008).
- ²²R. A. Fonseca, J. Vieira, F. Fiuza, A. Davidson, F. S. Tsung, W. B. Mori, and L. O. Silva, Plasma Physics and Controlled Fusion **55**, 124011 (2013).
- ²³J. Tonge, Ph.D. thesis, UCLA (2002).
- ²⁴V. Decyk, in *Int. Conf. on Plasma Physics (Kiev, USSR)* (Singapore: World Scientific, 1987).
- ²⁵T. Grismayer, J. E. Fahlen, V. K. Decyk, and W. B. Mori, Plasma Physics and Controlled Fusion **53**, 074011 (2011).
- ²⁶T. Takizuka and H. Abe, J. Comp. Phy. **25**, 205 (1977).
- ²⁷K. Nambu and S. Yonemura, J. Comp. Phy. **145**, 639 (1998).
- ²⁸F. Peano, M. Marti, L. O. Silva, and G. Coppa, Phys. Rev. E **79**, 025701 (2009).
- ²⁹T. N. Kato, arxiv:1312.5507 (2013).




FIG. 1. (Color online) Theoretical curves for the magnitude of the z component of the on axis electric field of a relativistic particle moving in the z direction. a) The strength of the electric field is dependent on cell size with larger cell sizes producing weaker fields. b) The strength of the electric field is also dependent on the interpolation order with higher order interpolation schemes producing weaker fields.




FIG. 2. (Color online) Reduction in stopping force due to finite-sized particles (“C” in eq. (19).) ‘Smoothing’ is used in finite-difference codes like Osiris, and here refers to a 4-pass (1,2,1) filter followed by a (-5,14,-5) compensator². ‘Filtering’ is done directly in k-space and therefore is used in spectral codes like Parsec; here the filter is $e^{\frac{-k^2 \Delta^2}{2}}$, where Δ is the cell size (applied to both the particles and the fields). From the figure it can be seen that smoothing is more effective at reducing stopping than higher-order particle shapes, and filtering is more effective than smoothing.




FIG. 3. (Color online) a) Wake formed in the E_1 field by a relativistic particle traveling along the x_1 direction in a 2D PIC simulation. This wake is predicted by wakefield theory. b) A comparison of the lineout along x_1 of the E_1 field near the relativistic particle predicted by theory (blue) and from the PIC simulation (red). The curves are in close agreement with the simulation curve (red) having a smaller amplitude and being slightly offset.




FIG. 4. (Color online) Relaxation of high-energy tails in a periodic simulation of a cold background and a hot electron tail. The blue curve shows the initial energy distribution, with the red, green and yellow curves showing the relaxation of the distribution with time, 92.75 fs apart when scaled to a density of $100n_c$. The curves show that the high-energy particles uniformly lose energy as predicted by theory.

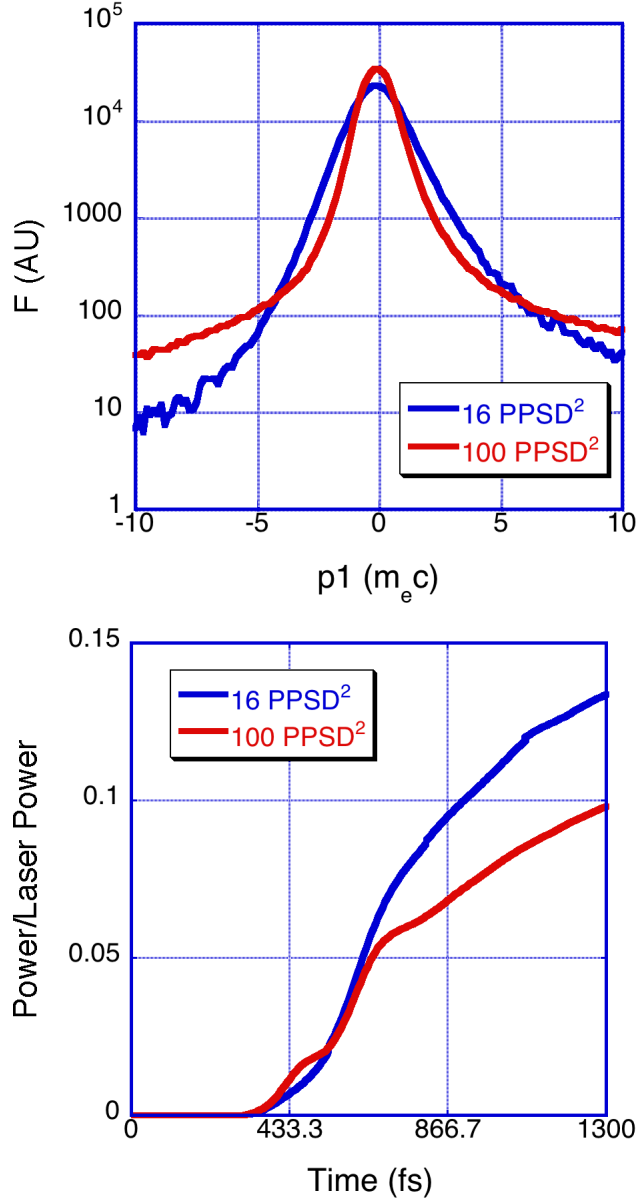


FIG. 5. (Color online) Panel a) shows the distribution function for forward momentum (p_1) $3 \mu\text{m}$ in front of the target core at 835 fs for 16PPSD² and 100PPSD² simulations respectively. These show that macro-particle stopping results in less heating of the plasma. Panel b) compares the power delivered to the core for the 16PPSD² and 100PPSD² simulations. These show that greater macro-particle stopping also results in lower energy electrons carrying energy to the target core.




FIG. 6. (Color online) Particle count vs. the logarithm of the kinetic energy (red), and the same data weighting each particle by the heatflux it carries in the x_1 direction (blue.) Arbitrary units, both independently normalized to a peak of 1.




FIG. 7. (Color online) Particle count as a function of γ and forward position for a) 2048 PPSD², b) 64 PPSD², and c) 64 PPSD² but with up to 5 particle splits for superthermal particles. d) Forward heatflux as a function of forward position for 2048 PPSD² (red curve), 64 PPSD² (magenta), 64 PPSD² w/5 splits (blue), and 64 PPSD² w/7 splits (pink). All data in arbitrary units but with equivalent normalization.